

A Multi-Year Urban Streetlight Imagery Dataset for Visual Monitoring and Spatio-Temporal Drift Detection

Peizheng Li^{a,*}, Ioannis Mavromatis^b, Ajith Sahadevan^a, Tim Farnham^a, Adnan Aijaz^a and Aftab Khan^a

^aBristol Research and Innovation Laboratory, Toshiba Europe Ltd., 30 Queen Square, Bristol, BS1 4ND, United Kingdom

^bDigital Catapult, London NW1 2RA, UK

ARTICLE INFO

Keywords:

Streetlight
Computer Vision
Machine Learning
Smart Cities
Drift Detection
Variational Autoencoders
Ablation
Multi-year Dataset

ABSTRACT

We present a large-scale, longitudinal visual dataset of urban streetlights captured by 22 fixed-angle cameras deployed across Bristol, U.K., from 2021 to 2025. The dataset contains over 526,000 images, collected hourly under diverse lighting, weather, and seasonal conditions. Each image is accompanied by rich metadata, including timestamps, GPS coordinates, and device identifiers. This unique real-world dataset enables detailed investigation of visual drift, anomaly detection, and MLOps strategies in smart city deployments. To promote secondary analysis, we additionally provide a self-supervised framework based on convolutional variational autoencoders (CNN-VAEs). Models are trained separately for each camera node and for day/night image sets. We define two per-sample drift metrics: relative centroid drift, capturing latent space deviation from a baseline quarter, and relative reconstruction error, measuring normalized image-domain degradation. This dataset provides a realistic, fine-grained benchmark for evaluating long-term model stability, drift-aware learning, and deployment-ready vision systems. The images and structured metadata are publicly released in JPEG and CSV formats, supporting reproducibility and downstream applications such as streetlight monitoring, weather inference, and urban scene understanding. The dataset can be found at <https://doi.org/10.5281/zenodo.17781192> and <https://doi.org/10.5281/zenodo.17859120>.

*Corresponding author

✉ peizheng.li@toshiba.eu (P. Li); ioannis.mavromatis@digicatapult.org.uk (I. Mavromatis); Ajith.Sahadevan@toshiba.eu (A. Sahadevan); Tim.Farnham@toshiba.eu (T. Farnham); adnan.aijaz@toshiba.eu (A. Aijaz); aftab.khan@toshiba.eu (A. Khan)

ORCID(s): <https://orcid.org/0000-0003-1516-1993> (P. Li); <https://orcid.org/0000-0002-3309-132X> (I. Mavromatis); <https://orcid.org/0009-0000-1046-7868> (A. Sahadevan); <https://orcid.org/0000-0002-5355-3982> (T. Farnham); <https://orcid.org/0000-0003-1048-0469> (A. Aijaz); <https://orcid.org/0000-0002-3573-6240> (A. Khan)

Specifications table

Subject	Computer Science
Specific subject area	Computer Vision, Drift Detection, Smart Cities Sensing
Type of data	Raw data in JPEG format and structured tabular data in CSV format. Each zipped folder contains image samples from a single lamppost camera (serial ID). The main CSV file includes metadata (timestamp, GPS, image name), photometric statistics, day/night status, and two drift metrics—relative centroid drift and relative reconstruction error. All metadata fields are provided in plain text, suitable for visual inspection, scripting, and downstream ML pipelines. Images are collected from 22 nodes and recorded over four years.
Data collection	Data were collected from 22 UMBRELLA nodes, each mounted on public streetlights in Bristol, UK. Each node includes a Raspberry Pi with Camera Module V1 (OV5647 sensor), capturing upward-facing images of the light column and sky. Images were collected hourly from 2021 to 2025 via the UMBRELLA testbed [1, 2]. Images were captured approximately once per hour, 24/7, for a period of four years (2021–2025). A randomised capture delay (up to 600 seconds) was added to reduce server load. All images were uploaded to a central server for indexing and post-processing.
Data source location	Around 7.2 km of public road in Bristol and South Gloucestershire, UK, including Coldharbour Lane, Filton Road, and the A4174 Ring Road. Additional nodes are located at the University of the West of England (Frenchay Campus).
Data accessibility	Data available online at the links below: https://doi.org/10.5281/zenodo.17781192 [3] https://doi.org/10.5281/zenodo.17859120 [4]
Related research article	None

1. Value of the Data

This work presents a large-scale, long-term visual dataset of streetlight imagery, purpose-built for research in urban infrastructure monitoring and Machine Learning-based drift detection. Its distinctive attributes make it valuable across a wide range of Computer Vision and Smart Cities applications. The key contributions are as follows:

- **Large-Scale, Longitudinal Visual Dataset:** The dataset comprises **526,096 JPEG images** collected between **2021 and 2025** from the city of Bristol, UK. Each image is captured by one of 22 fixed-angle Raspberry Pi Camera modules mounted on public lampposts. These devices provide diverse viewpoints with varying occlusions, scene compositions, and lighting conditions. The extended temporal span enables the analysis of both natural and operational distribution shifts, supporting tasks such as *drift detection*, *anomaly detection*, and long-term model evaluation.
- **Diverse Detection Scenarios:** Streetlights (and the nodes mounted on them) are manually classified into **Type-0 (ideal)** and **Type-1 (challenging)** categories, based on lamp visibility within the Field of View (FoV) of the camera. Type-0 nodes offer clear lamp visibility, enabling straightforward detection, whereas Type-1 nodes involve occlusions, misalignments, or poor lighting conditions. This distinction allows for robust benchmarking across varying degrees of difficulty.
- **Drift Detection Benchmark with Real-World Deployment:** We introduce a drift detection benchmark based on natural variations in urban scenes. The methodology includes: (1) quarterly temporal segmentation; (2)



Figure 1: (a) Type 0: Ideal lamppost in which the light is clearly visible; (b) Type 1: In this type, the light is not directly visible.

CNN-VAE training per day/night split and camera node; and (3) latent space evaluation using centroid drift and reconstruction error. An ablation across multiple latent dimensions quantifies sensitivity to representation capacity, enabling fair and reproducible comparisons.

- **Relevance to MLOps and Continual Learning:** The dataset reflects real-world visual evolution due to weather, lighting, and seasonal changes, providing a valuable testbed for validating MLOps strategies such as drift-triggered retraining, degradation monitoring, and lightweight continual learning models for edge devices. A number of potential use cases are described in [5].
- **Multimodal and Practical Applications:** Beyond drift detection, this dataset supports diverse Smart Cities applications, including:
 - Streetlight status recognition, classifying ON/OFF states for energy and maintenance optimisation;
 - Weather condition inference, e.g., fog or rain detection from image quality and noise;
 - Transfer learning for urban perception, due to the rich variations in lighting, occlusion, and environment.
- **Utility for Representation Learning:** The dataset’s variability in illumination, weather, and viewing angle makes it suitable for unsupervised representation learning. Pretraining on this data can yield robust, generalisable features for downstream Smart Cities vision tasks.

2. Background

The primary objective of this work is to introduce and open-source a large-scale, diverse, and temporally rich visual dataset of urban streetlights, captured via fixed-angle cameras installed on public lighting infrastructure across Bristol, UK. Spanning over **526,000 images** collected continuously between **2021 and 2025**, the dataset encompasses a wide spectrum of lighting conditions, seasonal changes, and weather-induced variations.

This dataset is designed to serve as a realistic and challenging benchmark for visual inspection, infrastructure monitoring, and maintenance research within Internet-of-Things (IoT) environments. Its uniqueness lies in its diverse FoV constraints and environmental artefacts such as rain, fog, and glare, which introduce complex edge cases ideal for advancing anomaly detection and robust classification under real-world deployment conditions.

To demonstrate its utility, we propose a drift detection framework based on Convolutional Variational Autoencoders (CNN-VAEs), leveraging two key metrics: latent space *centroid drift* and *relative reconstruction error*. These metrics are computed on a quarterly basis and provide fine-grained insight into evolving data distributions. Ablation studies across different latent dimensionalities and camera node types (e.g., visible vs. occluded lamp views) further explore the sensitivity of learned representations to environmental and architectural factors.

Overall, this dataset provides a valuable foundation for research in visual drift detection, unsupervised representation learning, model retraining policy development, and MLOps methodologies targeting long-term, real-world smart city deployments.

3. Data Description

The dataset is published across two Zenodo records [3, 4]. Record [3] contains node-specific archives for 18 streetlight nodes together with the central metadata file (streetcare-drift-dataset-2021-2025.csv), whereas record [4] includes the archives for the remaining four nodes. Each archive is a standalone .zip file that can be extracted and used independently.

Upon extraction, each archive yields a collection of JPEG images whose filenames are unique 32-character alphanumeric strings (e.g., f2a4b...c.jpg). These identifiers are randomly generated at capture time and remain stable across all processing stages. An example directory structure of the dataset is illustrated below:

```

streetcare-drift-dataset-2021-2025 [3, 4]
├── streetcare-drift-dataset-2021-2025.csv
├── RSE-19-C.zip
│   ├── 0a0e7888ad51010cfffdc14f20ede9564.jpg
│   ├── 00a0f83fda98ab9b586bee3a9335e4dd.jpg
│   └── ...
├── RSE-34-C.zip
│   ├── 0a0dbd9aace69a306bf73d36ed5daa4d.jpg
│   ├── 0a1c9bc590590e7536d43c4a5a898e95.jpg
│   └── ...
└── ...

```

All images were captured using the same Raspberry Pi Camera Module V1 [6], installed across 22 UMBRELLA nodes (i.e., lamppost-mounted devices). Each node offers a distinct viewpoint, due to variations in placement, height, orientation, and environmental occlusion. The dataset comprises a total of **526,096 images**, with **270,267 daytime** and **255,829 nighttime** captures, spanning a period of five years from 2021 to 2025. Images are further categorised into two visibility types (as shown in Fig. 1): Type-0 (**ideal nodes**, 258,527 images) and Type-1 (**challenging nodes**, 267,569 images), each encompassing data from 11 serial IDs, as summarised in Table 2.

Each UMBRELLA node captures an image approximately once per hour, with a randomised delay of up to 600 seconds to balance backend load. Captured images are transmitted to a backend server and indexed in a structured CSV file (each row in the CSV file represents an “entry” in the dataset), which includes metadata fields such as timestamp, serial ID, GPS coordinates, hostname, lighting condition (day/night), and fault status.

The CSV file also includes two important drift-related fields: `relative_centroid_drift` and `relative_recon_error`. These metrics enable per-sample drift analysis and are computed using a CNN-based Variational Autoencoder (CNN-VAE). Each image is resized to $64 \times 64 \times 3$ and passed through a trained CNN-VAE with a latent dimension of $d = 16$. The latent vectors \mathbf{z}_i are grouped quarterly by timestamp (`time_tag`) to track temporal drift. A baseline centroid is computed from the first quarter’s latent vectors.

- **Relative centroid drift** is calculated as the normalised Euclidean distance between each image’s latent vector and the baseline centroid.
- **Relative reconstruction error** measures the deviation of each image’s reconstruction loss from the baseline quarter’s average.

Both metrics allow granular, longitudinal quantification of distributional shift in both latent and image domains. They are particularly useful for MLOps-related applications [7] such as drift monitoring, retraining triggers, and continual learning policy evaluation. Users can recompute drift metrics using their own grouping strategy (e.g., monthly, biannual), since image-to-metadata alignment is maintained via the `image_name` field.

All data were collected through the UMBRELLA testbed described in [1], whose deployment topology is visualised in Fig. 2. The nodes are deployed along a 7.2 km urban corridor, covering diverse street and lighting conditions. A

Table 2
Dataset IDs, sizes, and number of images by the type of node

Type	Serial ID	Size (GB)	DayImgs	NightImgs	Type	Serial ID	Size (GB)	DayImgs	NightImgs
0	RSE-19-C	12	13,743	12,502	1	RSE-34-C	11	12,723	11,688
	RSS-A-2-C	9.9	11,026	10,858		RSE-44-C	12	13,361	12,695
	RSE-22-C	9.8	11,230	10,604		RSE-A-11-C	12	12,970	12,027
	RSE-21-C	9.7	12,133	11,501		RSE-A-15-C	12	13,550	12,379
	RSE-25-C	9.7	11,087	10,360		RSS-15-C	11	11,981	11,547
	RSE-24-C	9.0	12,631	11,905		RSS-47-C	11	11,722	11,340
	RSS-12-C	8.7	11,788	11,868		RSS-58-C	12	12,862	12,153
	RSE-6-C	8.3	12,968	11,707		RSS-59-C	11	11,956	11,552
	RSE-16-C	8.1	10,269	9,855		RSS-8-C	11	11,647	11,245
	RSE-A-10-C	8.1	12,328	12,124		RSS-A-13-C	11	12,094	11,599
	RSE-A-7-C	7.7	13,677	12,363		RSS-A-33-C	11	12,521	11,957

Table 3
Definition of the labels of each column of the CSV data files in the dataset.

Field Name	Definition
id	A unique ID given to each node. It increments by one for each image captured.
serial	The UMBRELLA node's "serial ID" (friendly name).
date	Date and timestamp at the moment an image was captured.
hostname	Hostname of the UMBRELLA node as seen in the OS.
lat	Latitude of the node location.
lon	Longitude of the node location.
image_name	JPEG image file name (32-digit alphanumeric string ending in '.jpg').
time_tag	Quarter when the image was captured (e.g., 2023-Q1).
fault_detected	The post-processed light status (operational or not).
confidence	The confidence score for the fault detection result.
daynight	Indicates whether the image was captured during the day or night cycle.
red	Mean pixel intensity in the red channel.
green	Mean pixel intensity in the green channel.
blue	Mean pixel intensity in the blue channel.
relative_centroid_drift	Normalized distance between a sample's latent vector and the baseline quarter's centroid.
relative_recon_error	Normalized reconstruction error relative to the baseline quarter's average.

representative comparison of the same lamppost over four years is illustrated in Fig. 3, showing natural degradation and increasing blur due to long-term wear, environmental exposure, and optical drift.

4. Drift Detection Experimental Methods and Results

4.1 Analysis Methods, Models, and Drift Metrics

To illustrate the utility of our dataset, we conduct a detailed camera drift analysis leveraging its long-term, high-frequency visual data. Given the extended temporal span and uncontrolled outdoor conditions, the dataset inherently exhibits non-stationary visual distributions. For our analysis, we segment the data by UMBRELLA node (using serial ID) and by lighting condition (day/night), with each image annotated by its timestamp and metadata.




Figure 2: The UMBRELLA network. Nodes are installed on public lampposts along a ~ 7.2 km road. Colours indicate connectivity: green = fibre, blue = WiFi, purple = fibre + LoRa gateway.

We adopt a self-supervised drift detection framework based on a CNN-VAE. A separate CNN-VAE model is trained for each node and lighting condition pair. All images are resized to $64 \times 64 \times 3$ before training. The encoder compresses input images into a latent vector $\mathbf{z} \in \mathbb{R}^d$, where d is the latent dimensionality. We conduct ablation over $d \in \{4, 8, 16, 32\}$. The decoder reconstructs the original image, and the model is trained using the standard VAE loss function: a combination of pixel-wise reconstruction error and KL divergence regularization. The model architecture is illustrated in Fig. 4.

To quantify drift, we compute latent and reconstruction-based metrics over quarterly temporal groups. As shown in Fig. 5, the latent space distributions for RSS-A-38-C (with $d = 32$) show clear separation between different quarters for both day and night images, suggesting a significant temporal shift. We define two complementary metrics:

- **Centroid Drift (CD):** For each quarter, we compute the mean latent vector (centroid) across samples. Drift is defined as the Euclidean distance between a given quarter’s centroid and the baseline centroid (from the first quarter).
- **Relative Reconstruction Error (RE):** For each quarter, we compute the mean pixel-wise reconstruction error and define drift as the relative increase from the baseline (first quarter’s) average.

Both metrics are globally normalized to enable comparisons across nodes, lighting conditions, and model configurations. Boxplots are used to visualize the distribution of these drift values across all 22 nodes, with separate analyses by latent dimension, day/night split, and visibility type (Type-0 vs. Type-1).

4.2 Results and Analysis

We trained a total of 44 CNN-VAE models, one for each of the 22 camera nodes, further split by day and night conditions. Each model was trained independently using 200 epochs, an Adam optimizer with a learning rate of 10^{-3} , and a batch size of 1024.

Post-training, we extracted latent vectors and reconstruction errors for all images. The data was temporally grouped into quarterly segments (i.e., 3-month intervals), yielding approximately 16 time bins for drift analysis. Results are visualized in Fig. 6 and Fig. 7, using boxplots to depict the distribution of normalized drift metrics across quarters. Each plot overlays results across four latent dimensions ($d \in \{4, 8, 16, 32\}$), separately for day/night splits and drift metrics.

- **Type-0 Nodes (Ideal FoV):**
 - Centroid drift remains consistently low, particularly for smaller latent dimensions and daytime images, suggesting that low-dimensional latent spaces are sufficient for modeling stable visual semantics under favorable conditions.
 - Larger latent dimensions (e.g., $d = 32$) introduce higher variance over time, possibly due to overfitting to minor visual changes or noise.
 - Reconstruction error exhibits a downward trend over time, indicating either lighting stabilization or reduced variability in captured scenes.
- **Type-1 Nodes (Challenging FoV):**




Figure 3: Example images captured by RSE-18-C in 2022 (a), 2023 (b), 2024 (c), 2025 (d).




Figure 4: VAE architecture for drift detection.

- Centroid drift increases steadily across quarters, particularly under night conditions. This trend likely reflects cumulative effects from occlusion, environmental change, or sensor aging.
- The increase in drift is more pronounced at higher latent dimensions, suggesting increased sensitivity to subtle or noisy variations.
- Reconstruction error varies non-monotonically, especially at night, likely due to inconsistent visibility or scene complexity.

These findings underscore the trade-off between model expressiveness and robustness. While larger latent spaces

A Dataset of Images of Public Streetlights for Monitoring and Drift Detection




Figure 5: Visualisation of the latent space distribution for the day dataset (left) and the night dataset (right) of serial ID RSE-19-C.




Figure 6: Boxplots of centroid drift and reconstruction error for ablation study results of type 0 serials.

can capture fine-grained changes, they are also more susceptible to overfitting and environmental noise. For long-term deployment scenarios such as Smart Cities monitoring or MLOps-driven retraining, a moderate latent dimension ($d = 8$ or $d = 16$) may offer the best balance between sensitivity and stability.

Overall, the results validate this dataset as a valuable benchmark for studying visual drift in real-world, long-term deployments. The proposed CNN-VAE framework, coupled with interpretable drift metrics, provides a scalable foundation for longitudinal model evaluation under realistic conditions.

A Dataset of Images of Public Streetlights for Monitoring and Drift Detection




Figure 7: Boxplots of centroid drift and reconstruction error for ablation study results of type 1 serials.

5. Limitations

The current release covers 22 upward-facing camera nodes deployed across a single UK municipality with identical hardware, mounting heights, and firmware. The annotations that accompany the images are intentionally lightweight: apart from node metadata, illumination regime, and automated quality-control tags, we do not provide exhaustive manual labels of lamp faults or seasonal events. This constrains the dataset to predominantly unsupervised or weakly supervised use cases, and practitioners seeking strongly labeled benchmarks for downstream detection or classification tasks will need to create their own annotations or fuse the imagery with municipal work-order records.

Ethics statements

Hereby, the authors consciously assure that for the manuscript “A Multi-Year Urban Streetlight Imagery Dataset for Visual Monitoring and Spatio-Temporal Drift Detection” the following is fulfilled:

1. This material is the authors’ own original work, which has not been previously published elsewhere.
2. The paper is not currently being considered for publication elsewhere.
3. The paper reflects the authors’ own research and analysis in a truthful and complete manner.
4. The paper properly credits the meaningful contributions of co-authors and co-researchers.
5. The results are appropriately placed in the context of prior and existing research.
6. All sources used are properly disclosed (correct citation). Literally copying of text must be indicated as such by using quotation marks and giving proper reference.
7. All authors have been personally and actively involved in substantial work leading to the paper, and will take public responsibility for its content.

As our dataset did not involve any human subjects, animal experiments, or social media platform data, approval from any IRB/local ethics committees was not required. As our camera images are facing the sky, no human subjects are present in the photos. Finally, as our dataset is based on street light images, no survey studies were conducted, and no work was conducted involving chemicals, procedures, or equipment that have any usual hazards inherent in their use, against animal or human subjects.

I agree with the above statements and declare that this submission follows the policies of Solid State Ionics as outlined in the Guide for Authors and in the Ethical Statement.

CRediT author statement

Peizheng Li: Conceptualization, Methodology, Software, Data Curation, Writing - Original Draft, **Ioannis Mavromatis:** Conceptualization, Software, Data Curation, Writing - Review, **Ajith Sahadevan:** Software, Data Curation, Writing - Review, **Tim Farnham:** Methodology, Writing - Review, **Adnan Aijaz:** Methodology, Writing - Review, **Aftab Khan:** Conceptualization, Methodology, Writing - Review & Editing.

Acknowledgments

This work is funded in part by Toshiba Europe Ltd. UMBRELLA project is funded in conjunction with South Gloucestershire Council by the West of England Local Enterprise Partnership through the Local Growth Fund, administered by the West of England Combined Authority.

Declaration of Competing Interest

- The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.
- The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

References

- [1] I. Mavromatis, Y. Jin, A. Stanoev, A. Portelli, I. Weeks, B. Holden, E. Glasspole, T. Farnham, A. Khan, U. Raza, A. Aijaz, T. Berton, I. Seto, N. Patel, and M. Sooriyabandara, "UMBRELLA: A One-stop Shop Bridging the Gap from Lab to Real-World IoT Experimentation," *IEEE Access*, Mar. 2024.
- [2] BRIL Toshiba Europe Ltd., "UMBRELLA Platform and Testbed," <https://www.umbrellaiot.com>, 2022, Accessed: 2022-1-31.
- [3] P. Li, I. Mavromatis, S. Ajith, T. Farnham, A. Aijaz, and A. Khan, "Multi-year urban streetlight image collection for visual monitoring and drift analysis (part 1/2)," Dec. 2025. [Online]. Available: <https://doi.org/10.5281/zenodo.17781192>
- [4] —, "Multi-year urban streetlight image collection for visual monitoring and drift analysis (part 2/2)," Dec. 2025. [Online]. Available: <https://doi.org/10.5281/zenodo.17859120>
- [5] P. Li, I. Mavromatis, and A. Khan, "Past, present, future: A comprehensive exploration of ai use cases in the umbrella iot testbed," in *2024 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*. IEEE, 2024, pp. 787–792.
- [6] Raspberry Pi, "Raspberry Pi Camera Module ver.1 Datasheet," https://www.arducam.com/downloads/modules/RaspberryPi_camera/RaspberryPi_Camera_Module_DS_V3.0.pdf, 2015, Accessed: 2022-1-31.
- [7] P. Li, I. Mavromatis, T. Farnham, A. Aijaz, and A. Khan, "Adapting mlops for diverse in-network intelligence in 6g era: Challenges and solutions," *arXiv preprint arXiv:2410.18793*, 2024.